

# On the Impossibility of a Quantum Sieve Algorithm for Graph Isomorphism

Cristopher Moore  
moore@cs.unm.edu  
University of New Mexico  
and the Santa Fe Institute

Alexander Russell  
acr@cse.uconn.edu  
University of Connecticut

February 1, 2008

## Abstract

It is known that any quantum algorithm for Graph Isomorphism that works within the framework of the hidden subgroup problem (HSP) must perform highly entangled measurements across  $\Omega(n \log n)$  coset states. One of the only known models for how such a measurement could be carried out efficiently is Kuperberg’s algorithm for the HSP in the dihedral group, in which quantum states are adaptively combined and measured according to the decomposition of tensor products into irreducible representations. This “quantum sieve” starts with coset states, and works its way down towards representations whose probabilities differ depending on, for example, whether the hidden subgroup is trivial or nontrivial.

In this paper we give strong evidence that no such approach can succeed for Graph Isomorphism. Specifically, we consider the natural reduction of Graph Isomorphism to the HSP over the wreath product  $S_n \wr \mathbb{Z}_2$ . We show, modulo a group-theoretic conjecture regarding the asymptotic characters of the symmetric group, that no matter what rule we use to adaptively combine quantum states, there is a constant  $b > 0$  such that no algorithm in this family can solve Graph Isomorphism in  $e^{b\sqrt{n}}$  time. In particular, such algorithms are essentially no better than the best known classical algorithms, whose running time is  $e^{O(\sqrt{n \log n})}$ .

## 1 Introduction

Peter Shor’s efficient quantum algorithms for order finding, factoring, and the discrete logarithm [Sho94], and Simon’s algorithm for determining the symmetry of a type of 2-1 function defined on  $\{0, 1\}^n$  [Sim94], led a frenzied charge to uncover the full algorithmic potential of a general purpose quantum computer. Creative invocations of the order-finding primitive yielded efficient quantum algorithms for a number of other number-theoretic problems [Hal02, Hal05]. As the field matured, these algorithms were roughly unified under the general framework of the *hidden subgroup problem*, where one must determine a subgroup  $H$  of a group  $G$  by querying an oracle  $f : G \rightarrow S$  known to have the property that  $f(g) = f(gh) \Leftrightarrow h \in H$ . Solutions to this general problem are the foundation for almost all known superpolynomial speedups offered by quantum algorithms over their classical counterparts (see [AJL06] for an important recent exception).

The algorithms of Simon and Shor essentially solve the hidden subgroup problem on abelian groups, namely  $\mathbb{Z}_2^n$  and  $\mathbb{Z}_n^*$  respectively. Since then, *non-abelian* hidden subgroup problems have received a great deal of attention (e.g. [HRTS00, GSVV01, FIM<sup>+</sup>03, MRS04, BCv05, HMR<sup>+</sup>06]). The cases where the ambient group  $G$  is the dihedral group or the symmetric group possess the most compelling algorithmic applications: in particular, we can reduce Graph Isomorphism for rigid graphs of size  $n$  to the case of the hidden subgroup problem over the symmetric group  $S_{2n}$ , or more specifically the wreath product  $S_n \wr \mathbb{Z}_2$ , where the hidden subgroup is promised to be either zero or of order two.

The basic approach to these problems, motivated by the community’s success with abelian groups, is

to use the oracle to prepare “coset states”  $\rho_H$  of the form

$$\rho_H = \frac{1}{|G|} \sum_c |cH\rangle \langle cH| ,$$

where  $|S\rangle$ , for a subset  $S \subset G$ , denotes the uniform superposition  $(1/\sqrt{|S|}) \sum_g |g\rangle$ . In the abelian case, one proceeds by computing the quantum Fourier transform of such coset states, measuring the resulting states, and appropriately interpreting the results. In the case of the symmetric group, however, determining  $H$  from a quantum measurement of coset states is far more difficult. In particular, no *product measurement* (that is, a measurement which treats each coset state independently) can efficiently determine a hidden subgroup over  $S_n$  [MRS05]; in fact, any successful measurement must be *entangled* over  $\Omega(n \log n)$  coset states at once [HMR<sup>+</sup>06].

Unfortunately, our species’ evolutionary history has given us little experience in thinking about highly-entangled quantum measurements. So far, there are only two known families of such measurements which we can actually perform efficiently on a quantum computer:

**Implementation of the Pretty Good Measurement.** Pioneered by Bacon, Childs, and van Dam, this approach seeks to efficiently implement an explicit measurement, the *pretty good measurement*. For some choices of groups and subgroups, this measurement is known to be optimal [BCvD05, BCv05, MR05b]; for arbitrary groups, it was recently shown to be sufficient when performed on  $\Omega(\log |G|)$  coset states [HKK06]. On a number of groups, such as the affine, Heisenberg, and dihedral groups, implementing the pretty good measurement corresponds to solving random cases of a combinatorial problem. However, it is not clear how to carry out this program for the symmetric groups, or even whether a corresponding combinatorial problem exists.

**Sieves generated by Clebsch-Gordan decomposition.** Kuperberg’s algorithm for the hidden subgroup problem in the dihedral group [Kup03] starts by generating a large number of coset states and subjecting each one to *weak Fourier sampling*, so that it lies inside a known irreducible representation. It then proceeds with an adaptive “sieve” process, at each step of which it judiciously selects pairs of states and measures them in a basis consistent with the *Clebsch-Gordan* decomposition of their tensor product into irreducible representations. This sieve continues until we obtain a state lying in an “informative” representation: namely, one from which information about the hidden subgroup can be easily extracted. We can visualize the sieve as a forest, where leaves consist of coset states, each internal node measures the tensor product of its parents, and the informative representations lie at the roots.

The sieve approach is especially attractive in cases like Graph Isomorphism, where all we need to know is whether the hidden subgroup is trivial or nontrivial. Specifically, suppose that the hidden subgroup  $H$  is promised to be either the trivial subgroup  $\{1\}$  or a *conjugate* of a known subgroup  $H_0$ . Assume further that there is an irreducible representation  $\sigma$  of  $G$  with the property that  $\sum_{h \in H_0} \sigma(h) = 0$ ; that is, a “missing harmonic” in the sense of [MR05a]. In this case, if  $H$  is nontrivial then the probability of observing  $\sigma$  under weak Fourier sampling of the coset state  $\rho_H$  is zero. More generally, the probability of observing  $\sigma$  when measuring any state which is *H-invariant*—that is, one which is fixed under right multiplication by any element of  $H$ —is zero. As we discuss below, the sieve process maintains the *H*-invariance of the observed states throughout its operation, so it follows that  $\sigma$  cannot appear at any time in the algorithm. If, on the other hand, one can guarantee that this sieve process *does* observe  $\sigma$  with significant probability when the hidden subgroup is trivial and the corresponding states are completely mixed, it gives us an algorithm to distinguish the two cases.

For example, if we consider the case of the hidden subgroup problem in the dihedral group  $D_n$  where  $H$  is either trivial or a conjugate of  $H_0 = \{1, m\}$  where  $m$  is an involution, then the sign representation  $\pi$  is a missing harmonic. Applying Kuperberg’s sieve, we observe  $\pi$  with significant probability after  $e^{O(\sqrt{n})}$  steps if  $H$  is trivial, while we can never observe it if  $H$  is of order 2. This framework has been successfully applied to groups of the form  $G^n$  by Alagić et al. [AMR06], suggesting that it might apply, *mutatis mutandis*, to a wide class of groups.

One difficulty in analyzing these sieve algorithms is that the algorithm designer has considerable flexibility in how she selects which pairs of representations to combine. We will show below, however, that modulo a representation-theoretic conjecture on the symmetric group, the hidden subgroup problem related to Graph Isomorphism cannot be solved efficiently by any algorithm in this family. Specifically,

no adaptive selection rule produces an efficient sieve whose transcript of observations statistically distinguishes the isomorphic and nonisomorphic cases, unless it combines  $e^{\Omega(\sqrt{n})}$  coset states. Since the best known classical algorithms for Graph Isomorphism run in time  $e^{O(\sqrt{n \log n})}$  [Bab80, BL83], quantum algorithms of this kind offer essentially no improvement over their classical counterparts.

The paper is organized as follows. We give a brief review of group representation theory in Section 2. In Section 3 we precisely describe the family of sieve algorithms outlined above. In Sections 4 and 5, we describe the probability distributions on the sequence of representations observed throughout a sieve algorithm. In Sections 6 and 7, we focus on wreath product groups of the form  $G \wr \mathbb{Z}_2$ , and show that the observed distributions completely fail to distinguish trivial and nontrivial subgroups unless we observe a special type of representation. In Section 8 we bound the probability of this event in terms of the characters of  $G$ . Then, in Section 9, we describe a conjectured bound on the characters of the symmetric group, and in Section 10 we show that this conjecture implies our main result. We conclude in Section 11 by proving that our conjecture holds in a number of cases.

## 2 Fourier analysis on finite groups

In this section we review the representation theory of finite groups. Our treatment is primarily for the purposes of setting down notation; we refer the reader to [Ser77] for a complete account.

Let  $G$  be a finite group. A *representation*  $\sigma$  of  $G$  is a homomorphism  $\sigma : G \rightarrow \mathbf{U}(V)$ , where  $V$  is a finite-dimensional Hilbert space and  $\mathbf{U}(V)$  is the group of unitary operators on  $V$ . The *dimension* of  $\sigma$ , denoted  $d_\sigma$ , is the dimension of the vector space  $V$ . Fixing a representation  $\sigma : G \rightarrow \mathbf{U}(V)$ , we say that a subspace  $W \subset V$  is *invariant* if  $\sigma(g) \cdot W = W$  for all  $g \in G$ . When  $\sigma$  has no invariant subspaces other than the trivial subspace  $\{0\}$  and  $V$  itself,  $\sigma$  is said to be *irreducible*.

If two representations  $\sigma$  and  $\sigma'$  are the same up to a unitary change of basis, we say that they are *equivalent*. It is a fact that any finite group  $G$  has a finite number of distinct irreducible representations up to equivalence and, for a group  $G$ , we let  $\widehat{G}$  denote a set of representations containing exactly one from each equivalence class. We often say that each  $\sigma \in \widehat{G}$  is the *name* of an irreducible representation, or an *irrep* for short.

The irreps of  $G$  give rise to the Fourier transform. Specifically, for a function  $f : G \rightarrow \mathbb{C}$  and an element  $\sigma \in \widehat{G}$ , define the *Fourier transform of  $f$  at  $\sigma$*  to be

$$\hat{f}(\sigma) = \sqrt{\frac{d_\sigma}{|G|}} \sum_{g \in G} f(g) \sigma(g) \ .$$

The leading coefficients are chosen to make the transform unitary, so that it preserves inner products:

$$\langle f_1, f_2 \rangle = \sum_g f_1^*(g) f_2(g) = \sum_{\sigma \in \widehat{G}} \text{tr} \left( \hat{f}_1(\sigma)^\dagger \cdot \hat{f}_2(\sigma) \right) \ .$$

If  $\sigma$  is *not* irreducible, it can be decomposed into a direct sum of irreps  $\tau_i$ , each of which acts on an invariant subspace, and we write  $\sigma \cong \tau_1 \oplus \cdots \oplus \tau_k$ . In general, a given  $\tau$  can appear multiple times in this decomposition, in the sense that  $\sigma$  may have an invariant subspace isomorphic to the direct sum of  $a_\tau$  copies of  $\tau$ . In this case  $a_\tau$  is called the *multiplicity* of  $\tau$  in the decomposition of  $\sigma$ .

There is a natural product operation on representations: if  $\lambda : G \rightarrow \mathbf{U}(V)$  and  $\mu : G \rightarrow \mathbf{U}(W)$  are representations of  $G$ , we may define a new representation  $\lambda \otimes \mu : G \rightarrow \mathbf{U}(V \otimes W)$  as  $(\lambda \otimes \mu)(g) : \mathbf{u} \otimes \mathbf{v} \mapsto \lambda(g)\mathbf{u} \otimes \mu(g)\mathbf{v}$ . This representation corresponds to the *diagonal action* of  $G$  on  $V \otimes W$ , in which we apply the same group element to both parts of the tensor product. In general, the representation  $\lambda \otimes \mu$  is not irreducible, even when both  $\lambda$  and  $\mu$  are. This leads to the *Clebsch-Gordan problem*, that of decomposing  $\lambda \otimes \mu$  into irreps.

Given a representation  $\sigma$  we define the *character* of  $\sigma$ , denoted  $\chi_\sigma$ , to be the trace  $\chi_\sigma(g) = \text{tr} \sigma(g)$ . As the trace of a linear operator is invariant under conjugation, characters are constant on the conjugacy classes of  $G$ . Characters are a powerful tool for reasoning about the decomposition of reducible representations. In particular, when  $\sigma = \bigoplus_i \tau_i$  we have  $\chi_\sigma = \sum_i \chi_{\tau_i}$  and, moreover, for  $\sigma, \tau \in \widehat{G}$ , we have the orthogonality conditions

$$\langle \chi_\sigma, \chi_\tau \rangle_G = \frac{1}{|G|} \sum_{g \in G} \chi_\sigma(g) \chi_\tau(g)^* = \begin{cases} 1 & \sigma = \tau \\ 0 & \sigma \neq \tau \end{cases} \ .$$

Therefore, given a representation  $\sigma$  and an irrep  $\tau$ , the multiplicity  $a_\tau$  with which  $\tau$  appears in the decomposition of  $\sigma$  is  $\langle \chi_\tau, \chi_\sigma \rangle_G$ . For example, since  $\chi_{\lambda \otimes \mu}(g) = \chi_\lambda(g) \cdot \chi_\mu(g)$ , the multiplicity of  $\tau$  in the Clebsch-Gordan decomposition of  $\lambda \otimes \mu$  is  $\langle \chi_\tau, \chi_\lambda \chi_\mu \rangle_G$ .

A representation  $\sigma$  is said to be *isotypic* if the irreducible factors appearing in the decomposition are all isomorphic, which is to say that there is a single nonzero  $a_\tau$  in the decomposition above. Any representation  $\sigma$  may be uniquely decomposed into maximal isotypic subspaces, one for each irrep  $\tau$  of  $G$ ; these subspaces are precisely those spanned by all copies of  $\tau$  in  $\sigma$ . In fact, for each  $\tau$  this subspace is the image of an explicit projection operator  $\Pi_\tau$  which can be written

$$\Pi_\tau = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(g) .$$

A useful fact is that  $\Pi_\tau$  commutes with the group action; that is, for any  $h \in G$  we have

$$\sigma(h) \Pi_\tau \sigma(h)^\dagger = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(hgh^{-1}) = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(h^{-1}gh)^* \sigma(g) = \frac{1}{|G|} \sum_{g \in G} d_\tau \chi_\tau(g)^* \sigma(g) = \Pi_\tau .$$

Our algorithms will perform measurements which project into these maximal isotypic subspaces and observe the resulting irrep name  $\tau$ . For the particular case of coset states, this measurement is called *weak Fourier sampling* in the literature; however, since we are interested in a more general process which in fact performs a kind of strong multiregister sampling on the original coset states, we will use the term *isotypic sampling* instead.

Finally, we discuss the structure of a specific representation, the (*right*) *regular* representation  $\mathbf{reg}$ , which plays an important role in the analysis below.  $\mathbf{reg}$  is given by the permutation action of  $G$  on itself. Specifically, let  $\mathbb{C}[G]$  be the *group algebra* of  $G$ ; this is the  $|G|$ -dimensional vector space of formal sums

$$\left\{ \sum_g \alpha_g \cdot g \mid \alpha_g \in \mathbb{C} \right\} .$$

(Note that  $\mathbb{C}[G]$  is precisely the Hilbert space of a single register containing a superposition of group elements.) Then  $\mathbf{reg}$  is the representation  $\mathbf{reg} : G \rightarrow \mathbf{U}(\mathbb{C}[G])$  given by linearly extending right multiplication,  $\mathbf{reg}(g) : h \mapsto hg$ . It is not hard to see that its character  $\chi_{\mathbf{reg}}$  is given by

$$\chi_{\mathbf{reg}}(g) = \begin{cases} |G| & g = 1 , \\ 0 & g \neq 1 , \end{cases}$$

in which case we have  $\langle \chi_{\mathbf{reg}}, \chi_\sigma \rangle_G = d_\sigma$  for each  $\sigma \in \widehat{G}$ . Thus  $\mathbf{reg}$  contains  $d_\sigma$  copies of each irrep  $\sigma \in \widehat{G}$ , and counting dimensions on each side of this decomposition implies

$$|G| = \sum_{\sigma \in \widehat{G}} d_\sigma^2 . \tag{1}$$

This equation suggests a natural probability distribution on  $\widehat{G}$ , the *Plancherel distribution*, which assigns to each irrep  $\sigma$  the probability  $\mathcal{P}_{\text{planch}}^G(\sigma) = d_\sigma^2/|G|$ . This is simply the dimensionwise fraction of  $\mathbb{C}[G]$  consisting of copies of  $\sigma$ ; indeed, if we perform isotypic sampling on the completely mixed state on  $\mathbb{C}[G]$ , or equivalently the coset state where the hidden subgroup is trivial, we observe exactly this distribution.

In general, we can consider subspaces of  $\mathbb{C}[G]$  that are invariant under left multiplication, right multiplication, or both; these subspaces are called *left*-, *right*-, or *bi-invariant* respectively. For each  $\sigma \in \widehat{G}$ , the maximal  $\sigma$ -isotypic subspace is a  $d_\sigma^2$ -dimensional bi-invariant subspace; it can be broken up further into  $d_\sigma$   $d_\sigma$ -dimensional left-invariant subspaces, or (transversely)  $d_\sigma$   $d_\sigma$ -dimensional right-invariant subspaces. However, this decomposition is not unique. If  $\sigma$  acts on a vector space  $V$ , then choosing an orthonormal basis for  $V$  allows us to view  $\sigma(g)$  as a  $d_\sigma \times d_\sigma$  matrix. Then  $\sigma$  acts on the  $d_\sigma^2$ -dimensional space of such matrices by left or right multiplication, and the columns and rows correspond to left- and right-invariant spaces respectively.

### 3 Clebsch-Gordan sieves

Consider the hidden subgroup problem over a group  $G$  with the added promise that the hidden subgroup  $H$  is either the trivial subgroup, or a conjugate of some fixed nontrivial subgroup  $H_0$ . We shall consider sieve algorithms for this problem that proceed as follows:

1. The oracle is used to generate  $\ell = \ell(n)$  coset states  $\rho_H$ , each of which is subjected to weak Fourier sampling. This results in a set of states  $\rho_i$ , where  $\rho_i$  is a mixed state known to lie in the maximal  $\sigma_i$ -isotypic subspace of  $\mathbb{C}[G]$  for some irrep  $\sigma_i$ .
2. The following *combine-and-measure* procedure is then repeated as many times as we like. Two states  $\rho_i$  and  $\rho_j$  in the set are selected according to an arbitrary adaptive rule that may depend on the entire history of the computation (in existing algorithms of this type, this selection in fact depends only on the irreps  $\sigma_i$  and  $\sigma_j$  in which they lie). We then perform isotypic sampling on their tensor product  $\rho_i \otimes \rho_j$ : that is, we apply a measurement operator which observes an irrep  $\sigma$  in the Clebsch-Gordan decomposition of  $\sigma_i \otimes \sigma_j$  (see [Kup03] or [MR05a] for how this measurement can actually be carried out by applying the diagonal action). This measurement destroys  $\rho_i$  and  $\rho_j$ , and results in a new mixed state  $\rho$  which lies in the maximal  $\sigma$ -isotypic subspace; we add this new state to the set.
3. Finally, depending on the sequence of observations resulting from the measurements above, the algorithm guesses the hidden subgroup.

We set down some notation to discuss the result of applying such an algorithm. Fixing a group  $G$  and a subgroup  $H$ , let  $A$  be a sieve algorithm which initially generates  $\ell$  coset states. As a bookkeeping tool, we will describe intermediate states of  $A$ 's progress as a *forest of labeled binary trees*. Throughout, we will maintain the invariant that the roots of the trees in this forest correspond to the current set of states available to the algorithm.

Initially, the state of the algorithm consists of a forest consisting of  $\ell$  single-node trees, each of which is labeled with the irrep name  $\sigma_i$  that resulted from performing weak Fourier sampling on a coset state, and is associated with the resulting state  $\rho_i$ . Then, each combine-and-measure step selects two root nodes,  $r_1$  and  $r_2$ , and applies isotypic sampling to the tensor product of their states. We associate the resulting state  $\rho$  with a new root node  $r$ , and place the nodes  $r_1$  and  $r_2$  below it as its children. We label this new node<sup>1</sup> with the irrep name  $\sigma$  observed in this measurement.

Thus, every node of the forest corresponds to a state that existed at some point during the algorithm, and each node  $i$  is labeled with the name of the irrep  $\sigma_i$  observed in the isotypic measurement performed when that node was created. We call the resulting labeled forest the *transcript* of the algorithm: note that this transcript contains all the information the algorithm may use to determine the hidden subgroup.

We make several observations about algorithms of this type. First, it is easy to see that nothing is gained by combining  $t > 2$  states at a time; we can simulate this with an algorithm which builds a binary tree with  $t$  leaves, and which ignores the results of all its measurements except the one at the root.

Second, the algorithm maintains the following kind of symmetry under the action of the subgroup  $H$ . Suppose we have a representation  $\sigma$  acting on a Hilbert space  $V$ . Given a subgroup  $H$ , we say that a state  $\psi \in V$  is *H-invariant* if  $\sigma(h) \cdot \psi = \psi$  for all  $h \in H$ . Similarly, given a mixed state  $\rho$ , we say that  $\rho$  is *H-invariant* if  $\sigma(h) \cdot \rho \cdot \sigma(h)^\dagger = \rho$  or, equivalently, if  $\sigma(h)$  and  $\rho$  commute. For instance, the coset state  $\rho_H$  is *H-invariant* under the right regular representation, since right-multiplying by any  $h \in H$  preserves each left coset  $cH$ . Now, suppose that  $\rho_1$  and  $\rho_2$  are *H-invariant*; clearly  $\rho_1 \otimes \rho_2$  is *H-invariant* under the diagonal action, and performing isotypic sampling preserves *H-invariance* since  $\Pi_\tau$  commutes with the action of any group element. Thus the states produced by the algorithm are *H-invariant* throughout.

Third, it is important to note that while at each stage we observe only an irrep name, rather than a basis vector inside that representation, by iterating this process the sieve algorithm actually performs a kind of *strong multiregister Fourier sampling* on the original set of coset states. For instance, in the dihedral group, suppose that performing weak Fourier sampling on two coset states results in the two-dimensional irreps  $\sigma_j$  and  $\sigma_k$ , and that we then observe the irrep  $\sigma_{j+k}$  under isotypic sampling of their tensor product. We now know that the original coset states were in fact confined to a particular subspace, spanned by two entangled pairs of basis vectors.

---

<sup>1</sup>New nodes is good nodes.

Finally, we note that the states produced by a sieve algorithm are quite different from coset states. In particular, they belong not to a maximal isotypic subspace of  $\mathbb{C}[G]$ , but to a (typically much higher-dimensional) *non-maximal* isotypic subspace of  $\mathbb{C}[G]^{\otimes \ell}$ , where  $\ell$  is the number of coset states feeding into that state (i.e., the number of leaves of the corresponding tree). Moreover, they have more symmetry than coset states, since each isotypic measurement implies a symmetry with respect to the diagonal action on the set of leaves descended from the corresponding internal node. In the next sections we will show how these states can be written in terms of projection operators applied to this high-dimensional space.

## 4 Observed distributions for fixed topologies

In general, the probability distributions arising from the combine-and-measure steps of a sieve algorithm depend on both the hidden subgroup and the entire history of previous measurements and observations (that is, the labeled forest, or transcript, describing the algorithm's history thus far). In this section and the next, we focus on the probability distribution induced by a *fixed forest topology* and subgroup  $H$ . We can think of this either as the probability distribution conditioned on the forest topology, or as the distribution of transcripts produced by some *non-adaptive* sieve algorithm, which chooses which states it will combine and measure ahead of time. We will show, modulo our conjecture, that for all forest topologies of sufficiently small size, the induced distributions on irrep labels fail to distinguish trivial and nontrivial subgroups. Then, in Section 10, we will complete the argument for adaptive algorithms.

Clearly, in this non-adaptive case the distributions of irrep labels associated with different trees in the forest are independent. Therefore, we can focus on the distribution of labels for a specific tree. At the leaves, the labels are independent and identically distributed according to the distribution resulting from weak Fourier sampling a coset state [HRTS00], namely

$$\mathcal{D}_H(\sigma) = \frac{d_\sigma |H|}{|G|} \mathbf{rk} \left( \sum_{h \in H} \sigma(h) \right) .$$

However, as we move inside the tree and condition on the irrep labels observed previously, the resulting distributions are quite different from this initial one. To calculate the resulting joint probability distribution, we need to define projection operators acting on  $\mathbb{C}[G]^{\otimes \ell}$  corresponding to the isotypic measurement at each node.

First, note that the coset state  $\rho_H$  can be written in the following convenient form:

$$\rho_H = \frac{1}{|G|} \sum_c |cH\rangle \langle cH| = \frac{1}{|G|} \sum_{h \in H} \mathbf{reg}(h)$$

where  $\mathbf{reg}$  is the *right regular representation* defined in Section 1: that is,  $\rho_H$  is proportional to the projection operator which right-multiplies by a random element of  $H$ ,

$$\Pi_H = \frac{1}{|H|} \sum_{h \in H} \mathbf{reg}(h) .$$

If  $H$  is trivial,  $\rho_H$  is the completely mixed state  $\rho_{\{1\}} = (1/|G|)\mathbf{1}$ . On the other hand, if  $H = \{1, m\}$  for an involution  $m$ , we have

$$\rho_H = \frac{1}{|G|} (1 + \mathbf{reg}(m))$$

which is proportional to the projection operator

$$\Pi_H = \frac{1}{2} (1 + \mathbf{reg}(m)) .$$

Now consider the tensor product of  $\ell$  “registers,” each containing a coset state. Given a linear operator  $M$  on  $\mathbb{C}[G]$  and a subset  $I \subseteq [\ell] = \{1, \dots, \ell\}$ , let  $M^I$  denote the operator on  $\mathbb{C}[G]^\ell \cong \mathbb{C}[G]^{\otimes \ell}$  which applies  $M$  to the registers in  $I$  and leaves the other registers unchanged. Then the mixed state consisting of  $\ell$  independent coset states can be written in a number of ways,

$$\rho_H^{\otimes \ell} = \frac{1}{|G|^\ell} \prod_{j=1}^{\ell} (1 + \mathbf{reg}(m)^{\{j\}}) = \frac{1}{|G|^\ell} \sum_{I \subseteq [\ell]} \mathbf{reg}(m)^I$$



and this state is proportional to a projection operator on  $\mathbb{C}[G^\ell]$ ,

$$\Pi_H^{\otimes \ell} = \frac{1}{2^\ell} \prod_{j=1}^{\ell} (1 + \mathbf{reg}(m)^{\{j\}}) = \frac{1}{2^\ell} \sum_{I \subseteq [\ell]} \mathbf{reg}(m)^I. \quad (2)$$

Note the sum over subsets of registers, a theme which has appeared repeatedly in discussions of multi-register Fourier sampling [BCvD05, HMR<sup>+</sup>06, Kup03, MR05a, MR05c, Reg02].

Now consider a tree  $T$  with  $\ell$  leaves corresponding to the  $\ell$  initial registers, and  $k$  nodes including the leaves. We can represent this tree as a set system, in which each node  $i$  is associated with the subset  $I_i \subseteq [\ell]$  of leaves descended from it. In particular,  $I_{\text{root}} = [\ell]$  and  $I_j = \{j\}$  for each leaf  $j$ .

Performing isotypic sampling at a node  $i$  corresponds to applying the diagonal action to its children (or in terms of the algorithm, its parents) and inductively to the registers in  $I_i$ : that is, we multiply each register in  $I_i$  by the same element  $g$  and leave the others fixed. If  $\sigma_i$  is the irrep label observed at that node, let us denote its character and dimension by  $\chi_i$  and  $d_i$  respectively, rather than the more cumbersome  $\chi_{\sigma_i}$  and  $d_{\sigma_i}$ . Then the projection operator corresponding to this observation is

$$\Pi_i^T = \frac{1}{|G|} \sum_{g \in G} d_i \chi_i(g)^* \mathbf{reg}(g)^{I_i} \quad (3)$$

Now consider a transcript of the sieve process which results in observing a set of irrep labels  $\sigma = \{\sigma_i\}$  on the internal nodes of the tree. The projection operator associated with this outcome is

$$\Pi^T[\sigma] = \prod_{i=1}^k \Pi_i^T. \quad (4)$$

We will abbreviate this as  $\Pi^T$  whenever the context is clear. Note that the various  $\Pi_i^T$  in the product (4) pairwise commute, since for any two nodes  $i, j$  either  $I_i$  and  $I_j$  are disjoint, or one is contained in the other. In the former case  $a^{I_i}$  and  $b^{I_j}$  for all  $a, b$ . In the latter case, say if  $I_i \subset I_j$ , we have  $a^{I_i} b^{I_j} = b^{I_j} (b^{-1} a b)^{I_i}$ , and since  $\chi_i(b^{-1} a b) = \chi_i(a)$  it follows from (3) that  $\Pi_i^T \Pi_j^T = \Pi_j^T \Pi_i^T$ .

Given a tree  $T$  with  $k$  nodes, we write  $P_T^{\{1\}}[\sigma]$  for the probability that we observe the set of irrep labels  $\sigma = \{\sigma_i\}$  in the case where the hidden subgroup is trivial. Since the tensor product of coset states is then the completely mixed state in  $\mathbb{C}[G^\ell]$ , this is simply the dimensionwise fraction of  $\mathbb{C}[G^\ell]$  consisting of the image of  $\Pi^T$ , or

$$P_T^{\{1\}}[\sigma] = \frac{1}{|G|^\ell} \text{tr} \Pi^T.$$

Moreover, since measuring a completely mixed state results in the completely mixed state in the observed subspace, each state produced by the algorithm is completely mixed in the image of  $\Pi^T$ . In particular, if the irrep label at the root of a tree is  $\sigma$ , the corresponding state consists of a classical mixture across some number of copies of  $\sigma$ , in each of which it is completely mixed. Thus, when combining two parent states with irrep labels  $\lambda$  and  $\mu$ , we observe each irrep  $\tau$  with probability equal to the dimensionwise fraction of  $\lambda \otimes \mu$  consisting of copies of  $\tau$ , namely

$$P_\tau^{\lambda \otimes \mu} = \frac{d_\tau}{d_\lambda d_\mu} \langle \chi_\tau, \chi_\lambda \chi_\mu \rangle_G \quad (5)$$

(recall that  $\langle \chi_\tau, \chi_\lambda \chi_\mu \rangle_G$  is the multiplicity of  $\tau$  in  $\lambda \otimes \mu$ ). We will refer to this as the *natural distribution* in  $\lambda \otimes \mu$ .

Now let us consider the case where the hidden subgroup is nontrivial. Since the mixed state  $\rho_{H^\ell}$  can be thought of as a pure state chosen randomly from the image of  $\Pi_H^{\otimes \ell}$ , the probability of observing a set of irrep labels  $\sigma$  in this case is

$$P_T^H[\sigma] = \frac{\text{tr} \Pi^T \Pi_H^{\otimes \ell}}{\text{tr} \Pi_H^{\otimes \ell}} = \frac{2^\ell}{|G|^\ell} \text{tr} \Pi^T \Pi_H^{\otimes \ell}$$

where we use the fact that  $\text{tr} \Pi_H^{\otimes \ell} = [G : H]^\ell = (|G|/2)^\ell$ . Below we abbreviate these distributions as  $P_T^{\{1\}}$  and  $P_T^H$  whenever the context is clear. Our goal is to show that, until the tree  $T$  is deep enough, these two distributions are extremely close, so that the algorithm fails to distinguish subgroups of the form  $\{1, m\}$  from the trivial subgroup.

## 5 Trivial and legal assignments

In this section we derive explicit expressions for  $P_T^{\{1\}}$  and  $P_T^H$ . First, we fix some additional notation. Given an assignment of group elements  $\{a_i\}$  to the nodes, for each leaf  $j$  we let  $\prod_{i \rightsquigarrow j} a_i$  denote the product of the elements along the path from the root to  $j$ :

$$\prod_{i \rightsquigarrow j} a_i = \prod_{i: j \in I_i} a_i$$

where the product is taken in order from the root to the leaf. Then using (3) and (4) we can write

$$\Pi^T = \frac{1}{|G|^k} \sum_{\{a_i\}} d_i \chi_i(a_i)^* \bigotimes_{j=1}^{\ell} \mathbf{reg} \left( \prod_{i \rightsquigarrow j} a_i \right)^{\{i\}}. \quad (6)$$

We say that an assignment  $\{a_i\}$  is *trivial* if  $\prod_{i \rightsquigarrow j} a_i = 1$  for every leaf  $j$ . Then, since  $\mathbf{tr} \mathbf{reg}(g) = \chi_{\mathbf{reg}}(g) = |G|$  if  $g = 1$  and 0 otherwise, we have

$$P_T^{\{1\}} = \frac{1}{|G|^\ell} \mathbf{tr} \Pi^T = \frac{1}{|G|^k} \sum_{\{a_i\} \text{ trivial}} \prod_{i=1}^k d_i \chi_i(a_i)^*. \quad (7)$$

To get a sense of how this expression scales, note that the particular trivial assignment where  $a_i = 1$  for all  $i$  contributes  $\prod_{i=1}^k d_i^2 / |G| = \prod_i \mathcal{P}_{\text{Planch}}(\sigma_i)$ , as if the  $\sigma_i$  were independent and Plancherel-distributed.

Now let us consider  $P_T^H$ . Combining (2) with (6) gives the following expression for  $\Pi^T \Pi_H^{\otimes \ell}$ :

$$\Pi^T \Pi_H^{\otimes \ell} = \frac{1}{2^\ell |G|^k} \sum_{\{a_i\}} d_i \chi_i(a_i)^* \bigotimes_{j=1}^{\ell} \mathbf{reg} \left( \left( \prod_{i \rightsquigarrow j} a_i \right) (1 + m) \right)^{\{i\}}. \quad (8)$$

We say that an assignment  $\{a_i\}$  is *legal* if  $\prod_{i \rightsquigarrow j} a_i \in \{1, m\}$  for every leaf  $j$ . Then the trace of the term corresponding to  $\{a_i\}$  is  $|G|^\ell$  if  $\{a_i\}$  is legal, and is 0 otherwise, and analogous to (7) we have

$$P_T^H = \frac{2^\ell}{|G|^\ell} \mathbf{tr} \Pi^T \Pi_H^{\otimes \ell} = \frac{1}{|G|^k} \sum_{\{a_i\} \text{ legal}} \prod_{i=1}^k d_i \chi_i(a_i)^*. \quad (9)$$

Thus these two distributions differ exactly by the terms corresponding to assignments which are legal but nontrivial. Our main result will depend on the fact that for most  $\sigma$  these terms are identically zero, in which case  $P_T^H$  and  $P_T^{\{1\}}$  coincide.

## 6 The wreath product and its representations

For any group  $G$ , the wreath product  $G \wr \mathbb{Z}_2$  is the semidirect product  $(G \times G) \rtimes \mathbb{Z}_2$ , where we extend  $G \times G$  by an involution which exchanges the two copies of  $G$ . Thus the elements  $((\alpha, \beta), 0)$  form a normal subgroup  $K \cong G \times G$  of index 2, and the elements  $((\alpha, \beta), 1)$  form its nontrivial coset. We will call these elements “non-flips” and “flips,” respectively.

The Graph Isomorphism problem reduces to the hidden subgroup problem on  $S_n \wr \mathbb{Z}_2$  in the following natural way. We consider the disjoint union of the two graphs, and consider permutations of their  $2n$  vertices. Then  $S_n \wr \mathbb{Z}_2$  is the subgroup of  $S_{2n}$  which either maps each graph onto itself (the non-flips) or exchanges the two graphs (the flips). We assume for simplicity that the graphs are rigid. Then if they are nonisomorphic, the hidden subgroup is trivial; if they are isomorphic,  $H = \{1, m\}$  where  $m$  is a flip of the form  $((\alpha, \alpha^{-1}), 1)$ , where  $\alpha$  is the permutation describing the isomorphism between them. Note that the set of such flips constitutes a conjugacy class.

For any group  $G$ , the irreps of  $G \wr \mathbb{Z}_2$  can be written in a simple way in terms of the irreps of  $G$ . It is useful to construct them by *inducing* upward from the irreps of  $K \cong G \times G$  (see [Ser77] for the definition of an induced representation). First, each irrep of  $K$  is the tensor product  $\lambda \otimes \mu$  of two irreps of  $G$ . Inducing this irrep from  $K$  up to  $G$  gives a representation

$$\sigma_{\{\lambda, \mu\}} = \text{Ind}_K^G(\lambda \otimes \mu)$$



of dimension  $2d_\lambda d_\mu$ . If  $\lambda \not\cong \mu$ , then this is irreducible, and  $\sigma_{\{\lambda, \mu\}} \cong \sigma_{\{\mu, \lambda\}}$  (hence the notation). We call these irreps *inhomogeneous*. Their characters are given by

$$\chi_{\{\lambda, \mu\}}((\alpha, \beta), t) = \begin{cases} \chi_\lambda(\alpha)\chi_\mu(\beta) + \chi_\mu(\alpha)\chi_\lambda(\beta) & \text{if } t = 0 \\ 0 & \text{if } t = 1 \end{cases} . \quad (10)$$

In particular, the character of an inhomogeneous irrep is zero at any flip.

On the other hand, if  $\lambda \cong \mu$ , then  $\sigma_{\{\lambda, \lambda\}}$  decomposes into two irreps of dimension  $d_\lambda^2$ , which we denote  $\sigma_{\{\lambda, \lambda\}}^+$  and  $\sigma_{\{\lambda, \lambda\}}^-$ . We call these irreps *homogeneous*. Their characters are given by

$$\chi_{\{\lambda, \lambda\}}^\pm((\alpha, \beta), t) = \begin{cases} \chi_\lambda(\alpha)\chi_\lambda(\beta) & \text{if } t = 0 \\ \pm \chi_\lambda(\alpha\beta) & \text{if } t = 1 \end{cases} . \quad (11)$$

In the next section, we will show that sieve algorithms obtain precisely zero information that distinguishes hidden subgroups of the form  $\{1, m\}$  from the trivial subgroup until it observes at least one homogeneous representation.

## 7 The importance of being homogeneous

Suppose that the irrep labels  $\sigma = \{\sigma_i\}$  observed during a run of the sieve algorithm consist entirely of inhomogeneous irreps of  $G \wr \mathbb{Z}_2$ . Since the irreps have zero character at any flip, the only trivial or legal assignments  $\{a_i\}$  that contribute to the sums (7) and (9) are those where each  $a_i$  is a non-flip, i.e., is contained in the subgroup  $K \cong G \times G$ . But the product of any string of such elements is also contained in  $K$ , so if this product is in  $H = \{1, m\}$  where  $m \notin K$ , it is equal to 1. Thus any legal assignment of this kind is trivial, the sums (7) and (9) coincide, and the probability of observing  $\sigma$  is the same in the trivial and nontrivial cases. That is, so long as every  $\sigma_i$  in  $\sigma$  is inhomogeneous,

$$P_T^H[\sigma] = P_T^{\{1\}}[\sigma] . \quad (12)$$

Our strategy will be to show that observing even a single homogeneous irrep is unlikely, unless the tree generated by the sieve algorithm is quite large. Moreover, because the two distributions coincide unless this occurs, it suffices to show that this is unlikely in the case where  $H$  is trivial.

Now, it is easy to see that the probability of observing a given representation in  $G \wr \mathbb{Z}_2$ , under either the Plancherel distribution or a natural distribution, factorizes neatly into the probabilities that we observe the corresponding pair of irreps, in either order, in a pair of similar experiments in  $G$ . First, the Plancherel measure of an inhomogeneous irrep  $\sigma_{\{\lambda, \mu\}}$  is

$$\mathcal{P}_{\text{planch}}^{G \wr \mathbb{Z}_2}(\sigma_{\{\lambda, \mu\}}) = \frac{(2d_\lambda d_\mu)^2}{2|G|^2} = 2 \mathcal{P}_{\text{planch}}^G(\lambda) \mathcal{P}_{\text{planch}}^G(\mu) . \quad (13)$$

Similarly, the probability that we observe a homogeneous irrep  $\sigma_{\{\lambda, \lambda\}}^\pm$  is the probability of observing  $\lambda$  twice under the Plancherel distribution in  $G$ , in which case the sign  $\pm$  is chosen uniformly:

$$\mathcal{P}_{\text{planch}}^{G \wr \mathbb{Z}_2}(\sigma_{\{\lambda, \lambda\}}^\pm) = \frac{d_\lambda^4}{|G|^2} = \mathcal{P}_{\text{planch}}^G(\lambda)^2 . \quad (14)$$

Now consider the natural distribution in the tensor product of two inhomogeneous irreps  $\sigma_{\{\lambda, \lambda'\}}$  and  $\sigma_{\{\mu, \mu'\}}$ . The multiplicity of a given homogeneous irrep  $\sigma_{\{\tau, \tau\}}^\pm$  in this tensor product factorizes as follows,

$$\left\langle \chi_{\{\tau, \tau\}}^\pm, \chi_{\{\lambda, \lambda'\}} \chi_{\{\mu, \mu'\}} \right\rangle_{G \wr \mathbb{Z}_2} = \frac{1}{2} \left( \langle \chi_\tau, \chi_\lambda \chi_{\lambda'} \rangle_G \langle \chi_\tau, \chi_\mu \chi_{\mu'} \rangle_G + \langle \chi_\tau, \chi_\lambda \chi_{\mu'} \rangle_G \langle \chi_\tau, \chi_{\lambda'} \chi_\mu \rangle_G \right) .$$

Thus the probability of observing either  $\sigma_{\{\tau, \tau\}}^+$  or  $\sigma_{\{\tau, \tau\}}^-$  under the natural distribution is

$$P_{\sigma_{\{\tau, \tau\}}^\pm}^{\sigma_{\{\lambda, \lambda'\}} \otimes \sigma_{\{\mu, \mu'\}}} = \frac{1}{2} \left( P_\tau^{\lambda \otimes \mu} P_\tau^{\lambda' \otimes \mu'} + P_\tau^{\lambda \otimes \mu'} P_\tau^{\lambda' \otimes \mu} \right) . \quad (15)$$

In other words, the probability of observing a homogeneous irrep of  $G \wr \mathbb{Z}_2$  is the probability of observing the same irrep in two natural distributions on  $G$ . Let us denote the probability that we observe the same irrep in the natural distributions in  $\lambda \otimes \mu$  and  $\lambda' \otimes \mu'$ —that is, that these two distributions collide—as

$$P_{\text{coll}}^{\lambda \otimes \mu, \lambda' \otimes \mu'} = \sum_{\tau} P_{\tau}^{\lambda \otimes \mu} P_{\tau}^{\lambda' \otimes \mu'} .$$

Then (15) implies that the total probability of observing a homogeneous irrep is

$$\sum_{\tau} P_{\sigma_{\{\tau, \tau\}}^{\pm}}^{\sigma_{\{\lambda, \lambda'\}} \otimes \sigma_{\{\mu, \mu'\}}} = \frac{1}{2} \left( P_{\text{coll}}^{\lambda \otimes \mu, \lambda' \otimes \mu'} + P_{\text{coll}}^{\lambda \otimes \mu', \lambda' \otimes \mu} \right) \leq \max \left( P_{\text{coll}}^{\lambda \otimes \mu, \lambda' \otimes \mu'}, P_{\text{coll}}^{\lambda \otimes \mu', \lambda' \otimes \mu} \right) . \quad (16)$$

In the next section, we will show that if  $\lambda, \mu, \lambda'$  and  $\mu'$  have certain character-theoretic properties, then no irrep  $\tau$  occurs too often in any of these natural distributions, and so the probability of a collision is small.

## 8 Collisions, smoothness, and characters

Let us bound the probability  $P_{\text{coll}} = P_{\text{coll}}^{\lambda \otimes \mu, \lambda' \otimes \mu'}$  that the natural distributions in  $\lambda \otimes \mu$  and  $\lambda' \otimes \mu'$  collide. The idea is that  $P_{\text{coll}}$  is small as long as both of either or both of these distributions is *smooth*, in the sense that they are spread fairly uniformly across many  $\tau$ . The following lemma shows that this notion of smoothness can be related to bounds on the normalized characters of typical representations.

**Lemma 1.** *Given a group  $G$ , say that an irrep  $\lambda$  of  $G$  is  $f(n)$ -smooth if*

$$\sum_{g \in G} \left| \frac{\chi_{\lambda}(g)}{d_{\lambda}} \right|^4 \leq f(n) .$$

*Suppose that  $\lambda$  and  $\mu$  are  $f(n)$ -smooth. Then*

$$P_{\text{coll}} \leq \frac{\max_{\tau} d_{\tau}}{\sqrt{|G|}} \sqrt{f(n)} .$$

*Proof.* We have

$$P_{\text{coll}} = \sum_{\tau} P_{\tau}^{\lambda \otimes \mu} P_{\tau}^{\lambda' \otimes \mu'} \leq \max_{\tau} P_{\tau}^{\lambda \otimes \mu} .$$

Using (5) and using Cauchy-Schwartz twice gives

$$\begin{aligned} P_{\tau}^{\lambda \otimes \mu} &= \frac{d_{\tau}}{d_{\lambda} d_{\mu}} \langle \chi_{\tau}, \chi_{\lambda} \chi_{\mu} \rangle_G \\ &\leq \frac{d_{\tau}}{d_{\lambda} d_{\mu}} \sqrt{\langle \chi_{\tau}, \chi_{\tau} \rangle_G \langle \chi_{\lambda} \chi_{\mu}, \chi_{\lambda} \chi_{\mu} \rangle_G} \\ &= \frac{d_{\tau}}{d_{\lambda} d_{\mu}} \sqrt{\langle |\chi_{\lambda}|^2, |\chi_{\mu}|^2 \rangle_G} \\ &\leq \frac{d_{\tau}}{d_{\lambda} d_{\mu}} \sqrt[4]{\langle |\chi_{\lambda}|^2, |\chi_{\lambda}|^2 \rangle_G \langle |\chi_{\mu}|^2, |\chi_{\mu}|^2 \rangle_G} \\ &= \frac{d_{\tau}}{\sqrt{|G|}} \sqrt[4]{\sum_{g \in G} \left| \frac{\chi_{\lambda}(g)}{d_{\lambda}} \right|^4 \sum_{g \in G} \left| \frac{\chi_{\mu}(g)}{d_{\mu}} \right|^4} \\ &\leq \frac{\max_{\tau} d_{\tau}}{\sqrt{|G|}} \sqrt{f(n)} . \end{aligned}$$

□

Now let us focus on the case relevant to Graph Isomorphism, where  $G = S_n$ . The reader can find an excellent introduction to the representation theory of  $S_n$  in [JK81]. Here we recall that each irrep of the

symmetric group  $S_n$  corresponds to a *Young diagram*, or equivalently an integer partition  $\lambda_1 \geq \lambda_2 \geq \dots$  where  $\sum_i \lambda_i = n$ . The number of irreps is thus the partition number,

$$p(n) = (1 + o(1)) \frac{1}{4\sqrt{3} \cdot n} e^{\delta\sqrt{n}} < e^{\delta\sqrt{n}} \quad (17)$$

where

$$\delta = \sqrt{2/3} \pi .$$

The maximum dimension of any irrep is bounded by the following result of Vershik and Kerov:

**Theorem 2 ([VK85]).** *There is a constant  $\hat{c} > 0$  such that for all  $n \geq 1$ ,*

$$\max_{\tau} d_{\tau} \leq e^{-(\hat{c}/2)\sqrt{n}} \sqrt{n!} .$$

In this case, Lemma 1 gives

$$P_{\text{coll}} \leq e^{-(\hat{c}/2)\sqrt{n}} \sqrt{f(n)} . \quad (18)$$

Therefore, our goal is to show that typical irreps of  $S_n$  are  $f(n)$ -smooth where  $f(n)$  grows slowly enough with  $n$ . Indeed, to prove our main result, even  $f(n) = e^{a\sqrt{n}}$  for some  $a < \hat{c}$  would suffice. However, a well-motivated conjecture, described in the next section, implies  $f(n)$ -smoothness with  $f(n) = O(\sqrt{n})$ .

## 9 A conjectured character bound

In this section we present a conjecture regarding the high-dimensional irreps of  $S_n$ . First we fix some notation. Given a permutation  $\pi \in S_n$ , let  $s(\pi)$  denote its *support*, i.e., the number of objects not fixed by it. Let  $c(\pi)$  denote the number of nontrivial cycles of which  $\pi$  is composed, i.e., other than fixed points, and let  $t(\pi)$  denote the length of the shortest string of transpositions whose product is  $\pi$ . Since a single  $k$ -cycle requires  $k - 1$  transpositions, we have  $t(\pi) = s(\pi) - c(\pi)$ .

Then the following conjecture states that for all irreps whose dimension is sufficiently close to the maximum dimension, their normalized characters obey a certain uniform bound.

**Conjecture 1.** *Say that an irrep  $\lambda$  of  $S_n$  is big if*

$$d_{\lambda} > e^{-\sqrt{n} \log n} \sqrt{n!} .$$

*Then there is a constant  $A$  such that, for  $n$  sufficiently large, the normalized character of all big  $\lambda$  obeys*

$$\left| \frac{\chi_{\lambda}(\pi)}{d_{\lambda}} \right| \leq A^{t(\pi)} n^{-t(\pi)/2}$$

*for all  $\pi \in S_n$ .*

We present evidence for this conjecture in Section 11 below. For now, we point out that it has the following two important consequences, which will imply our main result.

**Lemma 3.** *Conjecture 1 implies that all big irreps  $\lambda$  are  $O(\sqrt{n})$ -smooth.*

**Lemma 4.** *Say that an irrep  $\lambda$  of  $S_n$  is really big if*

$$d_{\lambda} > e^{-(1/2)\sqrt{n} \log n} \sqrt{n!} .$$

*Suppose that  $\lambda$  and  $\mu$  are really big, and that  $\tau$  is selected according to the natural distribution (5) in  $\lambda \otimes \mu$ . Then Conjecture 1 implies that  $\tau$  is really big with probability  $1 - e^{-\omega(\sqrt{n})}$ .*

We will prove Lemmas 3 and 4 with the help of the following combinatorial lemma.

**Lemma 5.** *Let  $0 < z < 1/n$ , and let*

$$q_n(z) = \sum_{\pi \in S_n} z^{t(\pi)} .$$

*Then*

$$q_n(z) = O(\sqrt{n}) \frac{e^{-n}}{(1 - zn)^{1/z}} . \quad (19)$$

*Proof.* Consider the exponential generating function

$$g(y, z) = \sum_{m=0}^{\infty} \frac{y^m}{m!} q_m(z) .$$

Using the techniques of [Wil94, Chapter 3], we can write this as a product over all  $k$  of contributions from the  $(k-1)!$  possible  $k$ -cycles, including fixed points. Since each such cycle contributes  $k$  to  $m$  and  $k-1$  to  $t(\pi)$ , and since there are  $(k-1)!$   $k$ -cycles on a given set of  $k$  objects, it follows that

$$g(y, z) = \prod_{k=1}^{\infty} \exp\left(\frac{y^k z^{k-1}}{k}\right) = \exp\left(\sum_{k=1}^{\infty} \frac{y^k z^{k-1}}{k}\right) = \exp\left(-\frac{1}{z} \ln(1-yz)\right) = \frac{1}{(1-yz)^{1/z}} .$$

Now note that  $e^{-n}g(n, z)$  is the expectation of  $q_m(z)$ , where  $m$  is Poisson distributed with mean  $n$ . Since  $q_m(z) > 0$ , this expectation is at least  $q_n(z)$  times the probability that  $m = n$ , which is  $e^{-n}n^n/n! = (1 - o(1))/\sqrt{2\pi n}$ . Thus we have

$$q_n(z) \leq (1 + o(1))\sqrt{2\pi n} \cdot e^{-n}g(n, z)$$

which concludes the proof.  $\square$

Note that since the Poisson distribution is tightly peaked, if  $q_n(z)$  varies smoothly with  $n$  then  $e^{-n}g(n, z)$  converges to  $q_n(z)$  as  $n \rightarrow \infty$ . In that case we could strengthen Lemma 5 by removing the  $O(\sqrt{n})$  term; but, since this is not necessary to our result, we have not attempted to prove it.

*Proof of Lemma 3.* If the conjecture holds, then we have

$$\sum_{\pi \in S_n} \left| \frac{\chi_{\lambda}(\pi)}{d_{\lambda}} \right|^4 \leq \sum_{\pi \in S_n} (A^{t(\pi)} n^{-t(\pi)/2})^4 = q_n(C/n^2) \quad (20)$$

where  $C = A^4$ . Setting  $z = C/n^2$  in Lemma 5 gives

$$q_n(C/n^2) = O(\sqrt{n}) e^{-n} \left(1 - \frac{C}{n}\right)^{-n^2/C} .$$

Using  $-\log(1-x) = x + x^2/2 + x^3/3 + \dots$ , we have

$$e^{-n} \left(1 - \frac{C}{n}\right)^{-n^2/C} = e^{-n} \exp\left(\left(\frac{C}{n} + O\left(\frac{1}{n^2}\right)\right) \left(\frac{n^2}{C}\right)\right) = e^{-n} e^{n+O(1)} = O(1)$$

which completes the proof.  $\square$

*Proof of Lemma 4.* Using simple dimension-counting arguments as in [MRS05], the multiplicity  $\langle \chi_{\tau}, \chi_{\lambda} \chi_{\mu} \rangle_G$  of  $\tau$  in  $\lambda \otimes \mu$  is at most  $d_{\tau}$ . First we show that we can assume that  $\tau$  is big; if it is not, by (5) we have

$$P_{\tau}^{\lambda \otimes \mu} \leq \frac{d_{\tau}^2}{d_{\lambda} d_{\mu}} < \frac{(e^{-\sqrt{n} \log n})^2}{(e^{-(1/2)\sqrt{n} \log n})^2} = e^{-\sqrt{n} \log n} = e^{-\omega(\sqrt{n})} . \quad (21)$$

Now, to bound the probability that  $\tau$  fails to be *really* big, we rewrite (5) as follows:

$$P_{\tau}^{\lambda \otimes \mu} = \frac{d_{\tau}^2}{n!} \sum_{\pi \in S_n} \frac{\chi_{\tau}}{d_{\tau}} \frac{\chi_{\lambda}}{d_{\lambda}} \frac{\chi_{\mu}}{d_{\mu}}$$

(recall that the characters of  $S_n$  are real). Since  $\tau$  is big, Conjecture 1 applies to it as well as to  $\lambda$  and  $\mu$ . In that case we can bound  $P_{\tau}^{\lambda \otimes \mu}$  using the triangle inequality,

$$P_{\tau}^{\lambda \otimes \mu} \leq \frac{d_{\tau}^2}{n!} \sum_{\pi \in S_n} \left| \frac{\chi_{\tau}}{d_{\tau}} \frac{\chi_{\lambda}}{d_{\lambda}} \frac{\chi_{\mu}}{d_{\mu}} \right| \leq \frac{d_{\tau}^2}{n!} \sum_{\pi \in S_n} (A^{t(\pi)} n^{-t(\pi)/2})^3 = \frac{d_{\tau}^2}{n!} q_n(C/n^{3/2})$$

where  $C = A^3$ . Setting  $z = C/n^{3/2}$  in Lemma 5 then gives

$$q_n(C/n^{3/2}) = O(\sqrt{n}) e^{-n} \left(1 - \frac{C}{n^{1/2}}\right)^{-n^{3/2}/C}.$$

Again using  $-\log(1-x) = x + x^2/2 + x^3/3 + \dots$ , we have

$$\begin{aligned} e^{-n} \left(1 - \frac{C}{n^{1/2}}\right)^{-n^{3/2}/C} &= e^{-n} \exp\left(\left(\frac{C}{n^{1/2}} + \frac{1}{2} \frac{C^2}{n} + O\left(\frac{1}{n^{3/2}}\right)\right) \left(\frac{n^{3/2}}{C}\right)\right) \\ &= e^{-n} e^{n+(C/2)\sqrt{n}+O(1)} = O(e^{(C/2)\sqrt{n}}). \end{aligned}$$

The probability that  $\tau$  fails to be really big, conditioned on the event that it is big, is thus less than

$$\frac{d_\tau^2}{n!} O(\sqrt{n}) O(e^{(C/2)\sqrt{n}}) e^{\delta\sqrt{n}} = e^{-\sqrt{n} \log n} e^{O(\sqrt{n})} = e^{-\omega(\sqrt{n})}. \quad (22)$$

Combining (21) and (22), the probability that  $\tau$  fails to be really big is  $e^{-\omega(\sqrt{n})}$ .  $\square$

## 10 Proof of the main result

We are now in a position to present our main result.

**Theorem 6.** *Let  $\hat{c}$  be the constant of Theorem 2. Then Conjecture 1 implies that for any constants  $b, c$  such that  $b < \hat{c}/2$  and  $c < (\hat{c}/2) - b$ , no sieve algorithm which combines less than  $e^{b\sqrt{n}}$  coset states can solve Graph Isomorphism with success probability greater than  $e^{-c\sqrt{n}}$ .*

*Proof.* Let us first consider the behavior of a sieve algorithm  $A$  in the case where the hidden subgroup  $H \subset S_n \wr \mathbb{Z}_2$  is trivial. For convenience, extend the notion of “really big” to representations of  $S_n \wr \mathbb{Z}_2$  by declaring  $\sigma_{\{\lambda, \mu\}}$  to be really big when both  $\lambda$  and  $\mu$  are. We will establish that with overwhelming probability, all the irrep labels observed by  $A$  are both really big and inhomogeneous.

Let  $\ell < e^{b\sqrt{n}}$  be the number of coset states initially generated by the algorithm. We begin by showing that with high probability, the irrep labels on the  $\ell$  leaves, i.e., those resulting from weak Fourier sampling, are all really big and homogeneous. If  $H$  is trivial, then these irrep labels are Plancherel-distributed, and by (13) the probability that a given one fails to be really big is at most twice the probability that a Plancherel-distributed irrep of  $S_n$  fails to be. The Plancherel measure of the set of non-really-big irreps of  $S_n$  is at most the number of irreps times the measure of a single such  $\lambda$ , so by (17) this probability is at most

$$2p(n) \frac{d_\lambda^2}{n!} < 2e^{\delta\sqrt{n}} e^{-\sqrt{n} \log n} = e^{-\omega(\sqrt{n})}.$$

Moreover, by (14) the probability that the label of a given leaf is homogeneous is the probability that we observe the same irrep of  $S_n$  in two samples of the Plancherel distribution. Using Theorem 2 this is

$$\sum_{\lambda} \left(\frac{d_\lambda^2}{|n!|}\right)^2 < \max_{\lambda} \frac{d_\lambda^2}{|n!|} \leq e^{-\hat{c}\sqrt{n}}.$$

Thus the combined probability that any of the  $\ell$  leaves have a label which is not both really big and inhomogeneous is

$$\ell \left(e^{-\omega(\sqrt{n})} + e^{-\hat{c}\sqrt{n}}\right) < e^{(b-\hat{c}+o(1))\sqrt{n}}. \quad (23)$$

Now, assume inductively that all the irreps observed by the algorithm before the  $i$ th combine-and-measure step are really big and inhomogeneous, and that the  $i$ th step combines states with two such labels  $\sigma_{\{\lambda, \lambda'\}}$  and  $\sigma_{\{\mu, \mu'\}}$ . Recall from (16) that the probability this results in a homogeneous irrep is bounded by the probability  $P_{\text{coll}}$  of a collision between a pair of natural distributions in  $S_n$ . Then Conjecture 1, Lemmas 1 and 3, and Theorem 2 imply that this probability is bounded by

$$P_{\text{coll}} \leq e^{-(\hat{c}/2)\sqrt{n}} O(n^{1/4}).$$

In addition, Conjecture 1 and Lemma 4 imply that the probability the observed irrep fails to be really big is  $e^{-\omega(\sqrt{n})}$ . Since each combine-and-measure step reduces the number of states by one, there are less than  $\ell$  such steps; taking a union bound over all of them, the probability that any of the observed irreps fail to be both homogeneous and really big is

$$\ell \left( e^{-(\hat{c}/2)\sqrt{n}} O(n^{1/4}) + e^{-\omega(\sqrt{n})} \right) < e^{(b-(\hat{c}/2)+o(1))\sqrt{n}}. \quad (24)$$

Let us call a transcript inhomogeneous if all of its irrep labels are. As  $c < (\hat{c}/2) - b$  by assumption, combining (23) and (24), for  $n$  sufficiently large,  $A$ 's transcript is inhomogeneous with probability at least  $1 - e^{-c\sqrt{n}}$ .

Now consider  $A$ 's behavior in the case of a nontrivial hidden subgroup,  $H = \{1, m\}$ . Inductively applying Equation (12) shows that the probability of observing any inhomogeneous transcript is exactly the same as it would have been if  $H$  were trivial. Thus the total variation distance between the distribution of transcripts generated by  $A$  in these two cases is no more than  $e^{-c\sqrt{n}}$ , and the theorem is proved.  $\square$

## 11 Evidence for the conjecture

In this section we give some evidence for our conjecture, by showing that it holds in a variety of cases. First we prove the following lemma, which bounds the height and width of a Young diagram corresponding to a big representation.

**Lemma 7.** *If  $\lambda$  is big, then the width and height of its Young diagram are less than  $4\sqrt{n} \log n$ .*

*Proof.* The Robinson-Schensted correspondence [Ful97] maps permutations to Young diagrams in such a way that the uniform measure on  $S_n$  maps to the Plancherel measure. In addition, the width (resp. height) of the Young diagram is equal to the length of the longest increasing (resp. decreasing) subsequence.

It follows that if  $\lambda$  has width  $w$ , its Plancherel measure  $d_\lambda^2/n!$  is at most the probability that a random permutation in  $S_n$  has an increasing subsequence of length  $w$ . By Markov's inequality, this probability is at most the expected number of such subsequences, which is

$$\binom{n}{w} \frac{1}{w!} < \frac{n^w}{w!^2} < \left( \frac{e\sqrt{n}}{w} \right)^{2w}$$

since  $w! > w^w e^{-w}$ . If  $w = 4\sqrt{n} \log n$ , comparing this to the Plancherel measure of a big irrep  $\lambda$  gives

$$e^{-2\sqrt{n} \log n} < \frac{d_\lambda^2}{n!} \leq \left( \frac{e}{4 \log n} \right)^{8\sqrt{n} \log n} < (4/e)^{-8\sqrt{n} \log n}$$

and since  $e < (4/e)^4$  this is a contradiction. The argument for the height is identical.  $\square$

Our next lemma uses deep results of Kerov [Ker03] and Biane [Bia98] to show that our conjecture holds for permutations whose support grows slowly enough with  $n$ . Indeed, it was Biane's work, which builds on recent results in the theory of noncommuting random variables, that originally inspired our conjecture.

**Lemma 8.** *Conjecture 1 holds for all permutations whose support is less than  $(\log n)^{1-\beta}$  for some  $\beta > 0$ .*

*Proof.* We first use a beautiful fact due to Kerov [Ker03]. Any sequence of irreps  $\lambda_n$  of  $S_n$  whose dimensions differ less than exponentially from the maximum dimension as  $n \rightarrow \infty$ , that is, such that  $d_{\lambda_n} = e^{-o(n)} \sqrt{n!}$ , have Young diagrams which converge to an asymptotic shape when scaled by  $\sqrt{n}$ . This condition is clearly satisfied by any sequence of big irreps; for aesthetic reasons, we show this asymptotic shape in Figure 1.

Note that these Young diagrams have width and height  $2\sqrt{n}$ , thus solving Ulam's celebrated problem of the likely length of the longest increasing or decreasing subsequence in a random permutation. However, the notion of convergence used in [Ker03] does not rule out the possibility that the Young diagrams corresponding to big  $\lambda$  have long thin "tails" near the horizontal and vertical axes, so we rely on Lemma 7 to bound the length of these tails.






Figure 1: The asymptotic shape of a Young diagram associated with an irrep  $\lambda$  of  $S_n$  whose dimension is close to the maximum dimension. It is scaled by  $\sqrt{n}$ , so its height and width are  $2\sqrt{n}$ .

Now, Theorem 1.3 of [Bia98] and the surrounding discussion establishes that for any sequence of irreps  $\lambda$  whose Young diagrams converge to some asymptotic shape, and which moreover have height and width  $o(n^{1/2+\epsilon})$  for all  $\epsilon > 0$  (which follows in our case from Lemma 7) the normalized character of any permutation with small support can be approximated by a product over its cycles. Specifically, for any asymptotic shape there is a set of constants  $a_k$  for  $k \geq 1$ , called the *free cumulants* associated with that shape,<sup>2</sup> such that, if  $s(\pi) < (\log n)^{1-\beta}$  for some fixed  $\beta > 0$ ,

$$\frac{\chi_\lambda(\pi)}{d_\lambda} = n^{-t(\pi)/2} \prod_{k \geq 1} a_k^{c_k(\pi)} + o(n^{-t(\pi)/2})$$

where  $c_k(\pi)$  is the number of cycles of size  $k$  in  $\pi$ . For the Kerov shape in particular,  $a_1 = 1$  and  $a_k = 0$  for all  $k \geq 2$ , giving a bound even tighter than our conjecture,

$$\frac{\chi_\lambda(\pi)}{d_\lambda} = o(n^{-t(\pi)/2})$$

for any  $\pi$  other than the identity. □

The next two lemmas use simple counting arguments to show that Conjecture 1 holds for permutations of large enough support, or which contain large cycles.

**Lemma 9.** *Conjecture 1 holds for all  $\pi$  with  $t(\pi) > \sqrt{n} \log n$ .*

*Proof.* We use the Murnaghan-Nakayama formula for the character [JK81]. A *ribbon tile* of length  $k$  is a polyomino of  $k$  cells, arranged in a path where each step is up or to the right. Given a Young diagram  $\lambda$  and a permutation  $\pi$  with cycle structure  $k_1 \geq k_2 \geq \dots$ , a *consistent tiling* consists of removing a ribbon tile of length  $k_1$  from the boundary of  $\lambda$ , then one of length  $k_2$ , and so on, with the requirement that the remaining part of  $\lambda$  is a Young diagram at each step. Let  $h_i$  denote the height of the ribbon tile corresponding to the  $i$ th cycle: then the Murnaghan-Nakayama formula states that

$$\chi_\lambda(\pi) = \sum_T \prod_i (-1)^{h_i+1} \quad (25)$$

where the sum is over all consistent tilings  $T$ .

Clearly the number of consistent tilings is an upper bound on  $|\chi_\lambda(\pi)|$ . Now, we claim that for any fixed  $k$ , the number of possible locations for a ribbon tile of length  $k$  on the boundary of a Young diagram  $\lambda$  of size  $n$  is less than  $\sqrt{2n}$ . To see this, associate each one with the cell of  $\lambda$  which is directly above

---

<sup>2</sup>Readers should be warned that our notation differs somewhat from Biane's.




Figure 2: We associate each possible location for a ribbon tile of fixed length  $k$  with a cell (shaded) which is above the tile's lower end and to the left of its upper end. The resulting sequence of cells moves up and to the right at each step, implying that the number of locations is less than  $\sqrt{2n}$ . Here  $k = 3$ .

the tile's lower end, and directly to the left of its upper end, as shown in Figure 2. A little thought reveals that the resulting sequence of cells has the property that each one is above and to the right of the previous one. Therefore, if there are  $\ell$  locations, we have

$$n \geq \sum_{i=1}^{\ell} i > \ell^2/2 .$$

and so  $\ell < \sqrt{2n}$ . It follows that the number of ways to remove the ribbon tiles corresponding to the  $c(\pi)$  nontrivial cycles is less than

$$(2n)^{c(\pi)/2} .$$

Moreover, after these ribbon tiles are removed, the number of consistent tilings of the remaining Young diagram is simply the dimension of the corresponding irrep of  $S_{n-s(\pi)}$ , which is less than  $\sqrt{|S_{n-s(\pi)}|} = \sqrt{(n-s(\pi))!}$ . Therefore, if  $\lambda$  is big we have

$$\begin{aligned} \left| \frac{\chi_{\lambda}(\pi)}{d_{\lambda}} \right| &< \frac{(2n)^{c(\pi)/2} \sqrt{(n-s(\pi))!}}{e^{-\sqrt{n} \log n} \sqrt{n!}} \\ &< 2 \cdot e^{\sqrt{n} \log n} 2^{c(\pi)/2} e^{s(\pi)/2} n^{(c(\pi)-s(\pi))/2} \\ &\leq 2 \cdot e^{\sqrt{n} \log n} (\sqrt{2}e)^{t(\pi)} n^{-t(\pi)/2} . \end{aligned}$$

Here we used the bound  $(n-s)!/n! < 4 \cdot n^{-s} e^s$ , implied by Stirling's approximation, and the facts that  $c(\pi) \leq t(\pi)$ ,  $s(\pi) \leq 2t(\pi)$ , and  $t(\pi) = s(\pi) - c(\pi)$ . Finally, if  $t(\pi) > \sqrt{n} \log n$ , the term  $e^{\sqrt{n} \log n}$  can be absorbed into  $A^{t(\pi)}$ , and Conjecture 1 holds for any  $A > \sqrt{2}e^2$ .  $\square$

The next lemma shows that Conjecture 1 also holds for permutations that contain sufficiently large cycles, for the simple reason that in this case the character is zero.

**Lemma 10.** *Conjecture 1 holds for all  $\pi$  that contain a cycle of size  $8\sqrt{n} \log n$  or greater.*

*Proof.* The maximum size of any ribbon tile that can be placed on the boundary of a Young diagram is its width plus its height (minus one). Thus if such a cycle exists, by Lemma 7 there are no consistent tilings and the Murnaghan-Nakayama formula (25) implies that the character is zero.  $\square$

Unfortunately, Lemmas 8, 9, and 10 leave uncovered the case of permutations whose support is between, roughly,  $\log n$  and  $\sqrt{n} \log n$  and which consist entirely of short cycles. However, it seems very plausible to us that the scaling we conjecture here extends across this gap.

Moreover, it should be noted our main result still holds even if the  $\log n$  term in our definition of big representations were replaced by a more slowly-growing function of  $n$  such as  $\log \log n$ . Thus we can get

by with significantly weaker forms of the conjecture, which apply only to representations even closer to the maximum dimension than we consider here.

## Acknowledgments

We are very grateful to Philippe Biane, Persi Diaconis, and Piotr Śniady for helpful conversations about the character theory of the symmetric groups, and to Sally Milius, Tracy Conrad and Rosemary Moore for inspiration. This work was supported by NSF grants CCR-0220070, EIA-0218563, and CCF-0524613, and ARO contract W911NF-04-R-0009.

## References

- [ACM06] ACM, editor. *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, New York, NY, USA, 2006. ACM Press.
- [AJL06] Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the Jones polynomial. In ACM [ACM06], pages 427–436.
- [AMR06] Gorjan Alagić, Cristopher Moore, and Alexander Russell. Subexponential-time algorithms for hidden subgroup problems over product groups. Technical Report quant-ph/0603251, arXiv.org e-Print archive, 2006.
- [Bab80] László Babai. On the complexity of canonical labeling of strongly regular graphs. *SIAM Journal on Computing*, 9(1):212–216, 1980.
- [BCv05] David Bacon, Andrew Childs, and Wim van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 469–478, 2005.
- [BCvD05] David Bacon, Andrew Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. Technical Report quant-ph/0501044, arXiv.org e-Print archive, 2005.
- [Bia98] Philippe Biane. Representations of symmetric groups and free probability. *Advances in Mathematics*, 138(1):126–181, 1998.
- [BL83] László Babai and Eugene M. Luks. Canonical labeling of graphs. In ACM, editor, *Proceedings of the fifteenth annual ACM Symposium on Theory of Computing, Boston, Massachusetts, April 25–27, 1983*, pages 171–183, New York, NY, USA, 1983. ACM Press. ACM order no. 508830.
- [FIM<sup>+</sup>03] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. In ACM, editor, *Proceedings of the 35th ACM Symposium on Theory of Computing*, pages 1–9, New York, NY, USA, 2003. ACM Press.
- [Ful97] William Fulton. *Young Tableaux: with Applications to Representation Theory and Geometry*, volume 35 of *Student Texts*. London Mathematical Society, 1997.
- [Gol94] Shafi Goldwasser, editor. *Proceedings: 35th Annual Symposium on Foundations of Computer Science, November 20–22, 1994, Santa Fe, New Mexico*, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1994. IEEE Computer Society Press. IEEE catalog number 94CH35717. IEEE Computer Society Press Order Number 6580-02.
- [GSVV01] Michelangelo Grigni, Leonard Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In ACM, editor, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 68–74, New York, NY, USA, 2001. ACM Press.
- [Hal02] Sean Hallgren. Polynomial-time quantum algorithms for Pell’s equation and the principal ideal problem. In ACM, editor, *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 653–658, New York, NY, USA, 2002. ACM Press.

- [Hal05] Sean Hallgren. Fast quantum algorithms for computing the unit group and class group of a number field. In ACM, editor, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 468–474, New York, NY, USA, 2005. ACM Press.
- [HMR<sup>+</sup>06] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. In ACM [ACM06], pages 604–617.
- [HRTS00] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. In ACM, editor, *Proceedings of the 32nd annual ACM Symposium on Theory of Computing: Portland*, pages 627–635, New York, NY, USA, 2000. ACM Press.
- [HKK06] Masahito Hayashi, Akinori Kawachi, and Hirotsada Kobayashi. General bounds for hidden subgroup problem. Technical Report quant-ph/0604174, arXiv.org e-Print archive, 2006.
- [JK81] Gordon James and Adalbert Kerber. *The representation theory of the symmetric group*, volume 16 of *Encyclopedia of mathematics and its applications*. Addison-Wesley, 1981.
- [Ker03] S. V. Kerov. *Asymptotic representation theory of the symmetric group and its applications in analysis*, volume 219 of *Translations of Mathematical Monographs*. American Mathematical Society, 2003. Translated by N. V. Tsilevich.
- [Kup03] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup. Technical Report quant-ph/0302112, arXiv.org e-Print archive, 2003.
- [MR05a] Cristopher Moore and Alexander Russell. Explicit multiregister measurements for hidden subgroup problems; or, Fourier sampling strikes back. Technical Report quant-ph/0504067, arXiv.org e-Print archive, 2005.
- [MR05b] Cristopher Moore and Alexander Russell. For distinguishing conjugate hidden subgroups, the Pretty Good Measurement is as good as it gets. Technical Report quant-ph/0501177, arXiv.org e-Print archive, 2005.
- [MR05c] Cristopher Moore and Alexander Russell. The symmetric group defies strong Fourier sampling: part II. Technical Report quant-ph/0501066, arXiv.org e-Print archive, 2005.
- [MRS04] Cristopher Moore, Alexander Russell, and Leonard Schulman. The power of basis selection in fourier sampling: hidden subgroup problems in affine groups. In *Proceedings of the fifteenth annual ACM-SIAM symposium on discrete algorithms*, pages 1113–1122, 2004.
- [MRS05] Cristopher Moore, Alexander Russell, and Leonard Schulman. The symmetric group defies Fourier sampling. In *Proceedings of the 46th Symposium on Foundations of Computer Science*, pages 479–488, 2005.
- [Reg02] O. Regev. Quantum computation and lattice problems. In IEEE, editor, *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2002, Vancouver, BC, Canada, 16–19 November 2002*, pages 520–529, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 2002. IEEE Computer Society Press. IEEE Computer Society Order Number PR01822.
- [Ser77] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Number 42 in Graduate Texts in Mathematics. Springer-Verlag, 1977.
- [Sho94] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In Goldwasser [Gol94], pages 124–134. IEEE catalog number 94CH35717. IEEE Computer Society Press Order Number 6580-02.
- [Sim94] D. R. Simon. On the power of quantum computation. In Goldwasser [Gol94], pages 116–123. IEEE catalog number 94CH35717. IEEE Computer Society Press Order Number 6580-02.
- [VK85] A. M. Vershik and S. V. Kerov. Asymptotic behavior of the maximum and generic dimensions of irreducible representations of the symmetric group. *Funk. Anal. i Prolizhen*, 19(1):25–36, 1985. English translation, *Funct. Anal. Appl.* 19(1):21–31, 1989.
- [Wil94] Herbert Wilf. *Generatingfunctionology*. Academic Press, 1994.